Resilient Supervisory Control of Autonomous Intersections in the Presence of Sensor Attacks

Amin Ghafouri and Xenofon D. Koutsoukos

arXiv:1606.06699v1 [cs.SY] 21 Jun 2016

***Abstract*— Cyber-physical systems (CPS), such as autonomous vehicles crossing an intersection, are vulnerable to cyber-attacks and their safety-critical nature makes them a target for malicious adversaries. This paper studies the problem of supervisory control of autonomous intersections in the presence of sensor attacks. Sensor attacks are performed when an adversary gains access to the transmission channel and corrupts the measurements before they are received by the decision-making unit. We show that the supervisory control system is vulnerable to sensor attacks that can cause collision or deadlock among vehicles. To improve the system resilience, we introduce a detector in the control architecture and focus on stealthy attacks that cannot be detected but are capable of compromising safety. We then present a resilient supervisory control system that is safe, non-deadlocking, and maximally permissive, despite the presence of disturbances, uncontrolled vehicles, and sensor attacks. Finally, we demonstrate how the resilient supervisor works by considering illustrative examples.**

I. Introduction

Cyber-physical systems (CPS) are vulnerable to cyber-attacks and their safety-critical nature makes them a target for malicious adversaries. In this paper, we study the problem of supervisory control of autonomous intersections in the presence of sensor attacks. Sensor attacks are performed when an adversary gains access to the transmission channel and corrupts the measurements before they are received by the decision-making unit. Such attacks can put the system at risk by interfering with the operation of the control system.

Recent advances in transportation technology indicate that autonomous intersections will be possible in the near future. The aim of autonomous intersections is to provide a safe, scalable, and efficient framework for coordinating autonomous vehicles. In the literature, there exist numerous protocols for autonomous intersections [12], [3], [14]. Among these approaches, studying the problem of autonomous intersection in the context of supervisory control of discrete event systems has received increasing attention since it allows incorporating the continuous dynamics and formally analyzing system safety [8], [10], [11], [1].

Supervisory control for autonomous intersections has three requirements: (1) safety, i.e., collisions must be avoided; (2) non-blockingness, i.e., vehicles should not deadlock; and (3) maximal-permissiveness, i.e., vehicles must not be restricted unless necessary [7]. A supervisory controller that satisfies these requirements is designed for a set of controlled and uncontrolled vehicles considering an unknown disturbance [10]. This is done by creating a finite discrete event system abstraction of the continuous system and formulating the problem in the context of supervisory control for discrete event systems. Further, the supervisor is reconstructed to handle bounded measurement uncertainties [11].

In this paper, we study the resilience of the system when deception attacks are performed on sensor measurements by a malicious omniscient adversary [6]. The problem is related to supervisory control with imperfect measurements [11]. However, sensor attack can be designed strategically and may have different effects from imperfect measurements.

Concerns about resiliency of supervisory control systems is not recent as it has been addressed in the area of fault-tolerant control [18], [22]. An approach for monitoring and diagnosis of hybrid systems that integrate model-based and statistical methods is proposed in [22]. Further, a framework for fault-tolerant control of discrete event systems is designed that ensures the system recovers from any fault [20]. However, the focus is on known faults rather than unknown and carefully designed cyber-attacks.

Cyber-security of control systems has been studied for different classes of systems. It is shown that knowledge of the physical system can be used for attack detection by considering malicious attacks on process control systems [5]. Denial-of-service attacks are studied for discrete-time linear dynamical systems [2]. Further, detection limitations for CPS are characterized in [16]. In particular, it is shown that an attack is undetectable if the measurements due to the attack coincide with the measurements due to some nominal operating condition.

Our objective is to study the problem of resilient supervisory control of autonomous intersections in the presence of sensor attacks. We introduce a detector in the control architecture with the purpose of detecting deception attacks on sensors. This detector incorporates knowledge of the physical system with the previously received data and detects a wide range of sensor attacks. Nevertheless, there exist stealthy attacks that cannot be detected but are capable of compromising safety. To address this issue, we design a resilient supervisor that maintains safety even in the presence of stealthy attacks. The resilient supervisor consists of an estimator that computes the smallest state estimate compatible with the control inputs and measurements seen thus far. We prove that the estimated set contains the true state of the system even if the measurements are corrupted and we show that the set of estimated states depends on the detector threshold. We then formulate the resilient supervisory control problem and present a solution. We demonstrate how the resilient supervisor works by considering illustrative examples.

The remainder of this paper is organized as follows. Section II defines the system model and sensor attacks, followed

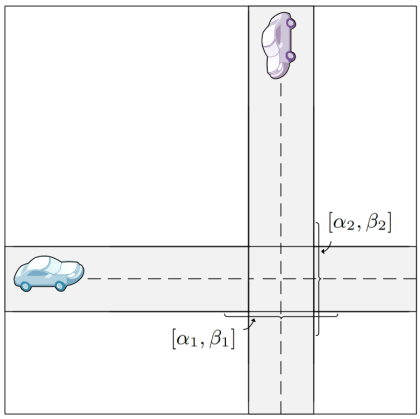

Fig. 1: Two-vehicle autonomous intersection

by formulation of the problem. In Section III, we introduce a detector in the system architecture and characterize stealthy attacks. We design the resilient supervisor in Section IV. In Section V, we illustrate how the resilient supervisor works by providing some examples. We conclude the paper in Section VI with discussion and future work.

II. System Model

A. Plant

We use the system model presented in [10]. Consider a set $\mathcal{N} = \{1, ..., n\}$ of vehicles driving towards an intersection. The vehicles are modeled as single integrators and their dynamics are described by

$$\dot{x} = v + d \tag{1}$$

where $x \in X \subset \mathbb{R}^n$ is the state (i.e., position), $v \in V \subset \mathbb{R}^n$ is the control input (i.e., speed), and $d \in D \subset \mathbb{R}^n$ is a disturbance input representing unmodeled dynamics. Assume $v \in V$ is a vector with elements in the finite set $\{\mu a, \mu(a+1), ..., \mu b\}$, where $\mu \in \mathbb{R}^+$ and $a, b \in \mathbb{N}$. We refer to $a\mu$ and $b\mu$ as v_{min} and v_{max}. The disturbance signal is bounded by $[d_{min}, d_{max}]^n$, and $v_{min} + d_{min} \geq \mu$, so that μ is the lowest possible speed. We assume that a subset of the vehicles are uncontrolled. To represent uncontrolled vehicles, we partition the vector v into two subvectors $v_c \in V_c$ and $v_{uc} \in V_{uc}$, such that $v = (v_c, v_{uc})$ and $V = V_c \times V_{uc}$, where v_c denotes the inputs of the controlled vehicles and v_{uc} denotes the inputs of the uncontrollable vehicles.

Assume the control input is kept constant over each time interval $[k\tau, (k+1)\tau]$. Define the time τ discretization of the system (1) as

$$x_{k+1} = x_k + u_k + \delta_k \tag{2}$$

where $x_k = x(k\tau)$, $u_k = v(k\tau)\tau$, and $\delta_k = \int_{k\tau}^{(k+1)\tau} d(t)dt$. Define $U = V\tau$ and $\Delta = D\tau$, and we have $u \in U$ and $\delta \in \Delta$. Also, let $U = U_c \times U_{uc}$, where U_c is the set of inputs for the controllable vehicles and U_{uc} is the set of inputs for the uncontrollable vehicles. Denote by $u = (u_c, u_{uc})$ the actions of the controllable and uncontrollable vehicles for any $u \in U$.

For each road l, let $[\alpha_l, \beta_l] \subset \mathbb{R}$ denote the part of the road that is in the intersection as shown in Fig. 1. Let $i \in \Pi_l$, if vehicle i drives along road l. A collision occurs between two vehicles $i \in \Pi_l$ and $j \in \Pi_{l'}$ with $l \neq l'$, whenever $x_i \in [\alpha_l, \beta_l]$ and $x_j \in [\alpha_{l'}, \beta_{l'}]$ simultaneously. Further, assuming $\gamma \in \mathbb{R}$ is the minimum safe distance between vehicles driving on the same road, a rear-end collision occurs when $|x_i - x_j| < \gamma$, $x_i \leq \beta_l$, $x_j \leq \beta_l$, and $i, j \in \Pi_l$. Calling the set of all the collision points the bad set $B \subset X$, no collision occurs if the system never reaches a state in B. A trajectory $x(t)$ of (1) is defined safe, if:

$$\inf_{t \geq 0, b' \in B} ||x(t) - b'||_\infty > 0 \tag{3}$$

B. Sensor Attacks and Problem Formulation

Sensor attacks are performed when an adversary compromises some sensors and corrupts the measurements before they are received by the decision-making unit [6]. Such attacks, which compromise integrity of the measurements, may lead to collision or deadlock among vehicles. To model sensor attacks, we assume $T_a = [k_s, k_e)$ is the period of attack and $e_i(k\tau) \in E \subseteq \mathbb{R}$, is an error added to sensor i at timestep $k \in T_a$. Defining $e(k\tau)$ as the vector $(e_1(k\tau), ..., e_n(k\tau))$, the corrupted measurement vector $\tilde{x}$ is described by

$$\tilde{x}(k\tau) = x(k\tau) + e(k\tau),\ k \in T_a \tag{4}$$

We assume there exists $T_{max} \in \mathbb{R}^+$ such that $|T_a| \leq T_{max}$. Because of the non-blockingness property, vehicles eventually cross the intersection, and therefore scenarios involving specific cars and intersections have a finite-time horizon. Since the sensor attack is launched on individual vehicles, the maximum attack duration, T_{max}, can be assumed to be finite. This variable can be selected as the maximum time it takes for a single vehicle to cross the intersection, which can easily be computed given the minimum speed (i.e., $v_{min} + d_{min}$) and the set of initial states. It should be noted that this implies an assumption that the initial state of a car is not corrupted. Although this assumption is restrictive, it is reasonable since detection of cars approaching intersections can be performed also with other means than the vehicles sensors. Also, if this assumption is not true an adversary can add "zombie" cars (i.e., cars that do not exist) or make existing cars "invisible", in which case, the detection and control problems considered in the paper are not well-defined.

Our goal is to design a supervisory controller that maintains the system's operational goals even in the presence of sensor attacks. These operational goals are: (1) safety, i.e., collisions must be avoided; (2) non-blockingness, i.e., vehicles must eventually cross the intersection; and (3) maximal-permissiveness, i.e., vehicles must not be restricted unless necessary [10].

III. Attack Detection

To be resilient to sensor attacks, a common approach is to employ detection mechanisms that detect attacks before they can cause significant damage.

A. Detector

We use a nonparametric Cumulative sum (CUSUM) statistic as our detection method [4]. The detector's goal is to determine whether a sequence of vectors $z(1), z(2), ..., z(N)$ is generated through the system being under normal behavior (H_0) or attack (H_1). Similar to [5], to formulate the detector and to avoid making limiting assumptions on the attacker's behavior, we assume the expected value of a random process $Z(k)$ that generates $z(k)$ is less than zero in the case of normal behavior, and greater than zero in the case of attack.

We denote by $\mathbf{Post}_{u_c} x$, the set of all the reachable states given a state x and an input u_c, that is, $\mathbf{Post}_{u_c}(x) := \cup_{u_{uc} \in U_{uc}, d \in D} (x + u + d)$. Given a measurement $\tilde{x}_i(k)$, the parameter $z_i(k)$ is defined for every vehicle i as

$$z_i(k) := \inf_{\hat{x}_i(k) \in \mathbf{Post}_{u_i}(\tilde{x}_i(k-1))} ||\tilde{x}_i(k) - \hat{x}_i(k)|| - b_i \quad (5)$$

where $b_i \in \mathbb{R}$ is a small constant chosen in a way that $E_{H_0}[z_i(k)] < 0$.

The CUSUM statistic is described by:

$$C_i(k) = (C_i(k-1) + z_i(k))^+ \quad (6)$$

where $C_i(0) = 0$, and $(a)^+ = a$ if $a \geq 0$.

Assigning η_i as the threshold chosen based on a desired false alarm rate, the corresponding decision rule is defined as

$$d(C_i(k)) = \begin{cases} H_1 & \text{if } C_i(k) > \eta_i \\ H_0 & \text{otherwise} \end{cases} \quad (7)$$

We assume that upon detection (i.e., $d(C_i(k)) = H_1$ for a k), the system switches to a mode where vehicles are either driven manually or controlled by a fail-safe controller.

B. Stealthy Attacks

Although the detector can identify attacks that cause a large difference between received measurements and predicted states (i.e., detectable attacks), a skilled attacker may attempt to evade this scheme via performing *stealthy attacks* [5]. This is done by keeping the detector's output below the threshold of detection and not triggering an alert. Stealthy attacks exist in our system because of the following factors:
(i) Detector's threshold: As shown in Fig. 2, small thresholds (i.e., $\eta \to 0$) reduce the size of stealthy attacks space. This means a system designer should always select the smallest possible threshold to achieve maximum resiliency. However, employment of detectors with small thresholds results in high false alarm rates in realistic noisy environments. This is why thresholds are often chosen non-zero (i.e., $\eta \neq 0$), which allows an adversary to launch stealthy attacks.
(ii) Disturbance and uncontrolled vehicles: Disturbance and uncontrolled vehicles add non-deterministic behavior to the system. This makes it impossible to perform exact state estimation. Hence, as long as the measurements due to the attack cannot be differentiated from the measurements due to some normal operation, the attack remains stealthy [16].

Stealthy attacks must be characterized and countermeasures for securing the system against them must be designed. To do so, we characterize stealthy attacks for the detector

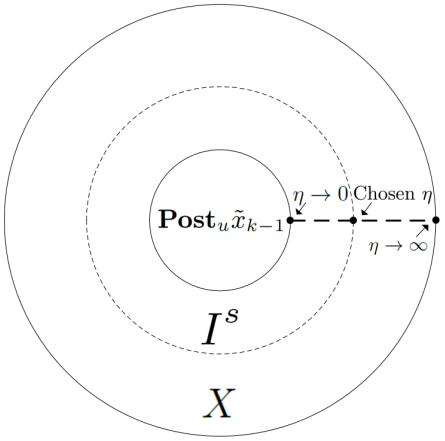

Fig. 2: Choosing a small threshold for the detector reduces size of stealthy attack space I^s.

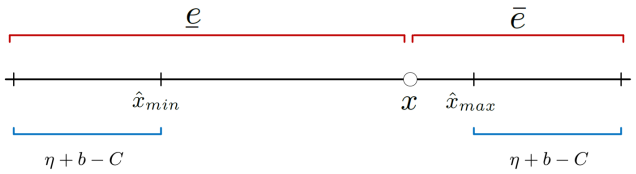

Fig. 3: Sensor attack is stealthy if $\tilde{x} \in [\hat{x}_{min} - (\eta + b - C), \hat{x}_{max} + \eta + b - C]$.

(7). We do this by deriving a condition for error vector e_k given the actual state $x(k\tau)$, which is only known by the attacker. We define $\hat{x}_{max}(k\tau) := \max \mathbf{Post}_{u_{k-1}}(\tilde{x}_{k-1})$ and $\hat{x}_{min}(k\tau) := \min \mathbf{Post}_{u_{k-1}}(\tilde{x}_{k-1})$ [1].

Lemma 1: (Stealthy Attacks): A sensor attack with $e_T \in E^{n \times |T|}$ is stealthy iff for all $k \in T_a$ we have $e_k \in [\underline{e}_k, \bar{e}_k]$ where

$$\underline{e}_k(x(k), C(k-1)) = -x(k) + \hat{x}_{min}(k) - \eta - b + C(k-1),$$

$$\bar{e}_k(x(k), C(k-1)) = -x(k) + \hat{x}_{max}(k) + \eta + b - C(k-1).$$

Proof: We show a corrupted vector $\tilde{x}(k)$ bypasses the detector if and only if it is within the stated bound. Let $\tilde{x}(k) \in X$ be a stealthy attack $\Leftrightarrow C(k) < \eta \Leftrightarrow C(k-1) + z(k) < \eta \Leftrightarrow C(k-1) + \inf ||\tilde{x}(k) - \hat{x}(k)|| - b < \eta \Leftrightarrow$ for each i we have $\inf |\tilde{x}_i(k) - \hat{x}_i(k)| < \eta_i + b_i - C_i(k-1) \Leftrightarrow \inf |x_i(k) + e_i(k) - \hat{x}_i(k)| < \eta_i + b_i - C_i(k-1)$. Solving this, we have $-x_i(k) + \hat{x}_{i,min}(k) - \eta_i - b_i + C_i(k-1) < e_i(k) < -x_i(k) + \hat{x}_{i,max}(k) + \eta_i + b_i - C_i(k-1) \Leftrightarrow \underline{e} < e_k < \bar{e}$. ■

The above lemma characterizes stealthy attacks by deriving a condition for error vector e_k assuming the actual state is known. We use (4) and Lemma 1 to define a set of stealthy attacks $I_k^s \subseteq X$, which contains all the corrupted measurements $\tilde{x}_k$ that are stealthy.

Lemma 2: (Stealthy Attacks Space): A sensor attack with $e_T \in E^{n \times |T|}$ is undetectable iff for all $k \in T_a$, the measurement $\tilde{x}_k$ belongs to stealthy attacks set $I_k^s : X \times U_c \times \mathbb{R} \to 2^X$

[1] We temporarily abuse notation and exclude τ when referring to a time $k\tau$.

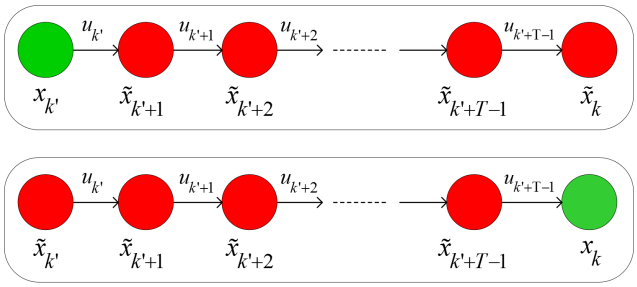

Fig. 4: Assuming $k = k' + T_{max}$, either $\tilde{x}_{k'}$ or $\tilde{x}_k$ is not attacked.

described by

$$I_k^s(\tilde{x}_{k-1}, u_{k-1}, C_{k-1}) = [\hat{x}_{min,k} - \eta - b + C, \hat{x}_{max,k} + \eta + b - C] \quad (8)$$

IV. Resilient Supervisory Control System

In this section, we design a resilient supervisory control system that is secure against stealthy attacks. This is done by taking the following steps: (1) Designing an estimator system that computes the smallest state estimate compatible with the sequence of control inputs and possibly corrupted measurements received; (2) Formulating the supervisory control problem given the estimator system with respect to a state space discretization; (3) Solving the problem by creating a discrete event system (DES) abstraction of the estimator system, translating the control problem to the abstracted domain, solving it, and translating the results back to the continuous domain.

A. Estimator

We present a method to compute the smallest state estimate compatible with the sequence of control inputs and possibly corrupted measurements received. We first use the assumption on maximum duration of sensor attacks to find a set that surely contains the actual state of the system even if measurements are corrupted. We then present a prediction-correction scheme based on the work in [11], which computes a set of predicted states and then corrects it. Finally, we design an estimator system that receives as inputs the measurements, the control inputs, and the detector's statistics and computes the smallest state set that contains the actual state.

In the following, we assume $k' = k - T_{max}$ and we define an extension of post-operator as $\mathbf{Post}_{u_{c,1}\dots u_{c,m}} x = x + \sum_{i=1}^{i=m} u_{c,i} + \bigcup_{u_{uc,i}\in U_{uc}, \delta_i \in \Delta} \sum_{i=1}^{i=m} u_{uc,i} + \delta_i$.

Lemma 3: Given measurements $\tilde{x}_k$ and $\tilde{x}_{k'}$, and previously chosen controllable inputs $(u_{c,k'}, ..., u_{c,k-1})$, output of function $\hat{I}_k : X \times \underbrace{U_c \ \dots \ U_c}_{T_{max}} \times X \to 2^X$ defined by

$$\hat{I}_k(\tilde{x}_{k'}, u_{k'}, ..., u_k, \tilde{x}_k) = \{\textbf{\textit{Post}}_{u_{c,k'}\dots u_{c,k-1}} \tilde{x}_{k'}\} \cup \tilde{x}_k \quad (9)$$

contains the actual state of the system for any timestep $k \in \mathbb{N}$, that is, $x_k \in \hat{I}_k, \ \forall k \in \mathbb{N}$.

Proof: We directly use the assumption that length of attacks is at most T_{max}. As shown in Fig. 4, this implies that $\tilde{x}_{k'}$ and $\tilde{x}_k$ may not be corrupted at the same time. Let us consider both cases. If $\tilde{x}_k$ is corrupted, T_{max} steps before (i.e., $k - T_{max} = k'$), the received measurement is uncorrupted, that is to say, $\tilde{x}_{k'} = x_{k'}$. Hence, we can write $x_k \in \mathbf{Post}_{u_{k'}\dots u_{k-1}} x_{k'} \subseteq \hat{I}_k$. Next, suppose $\tilde{x}_{k'}$ is corrupted, which means the measurement received at timestep $k = k' + T_{max}$ is accurate, i.e., $\tilde{x}_k = x_k$. This is followed by $x_k \in \hat{I}_k$. Thus, $\hat{I}_k$ contains the true state of the system in both cases. ■

Lemma 3 holds for any two measurements that have a time distance of at least T_{max}. This gives rise to definition of the prediction-correction scheme below that recursively computes set of predicted states $I^p \subseteq 2^X$ and set of corrected states $I^c \subseteq 2^X$. We assume the initial state set X_0 is not in the bad region. Also, for $k < T_{max}$, we define $\hat{I}_k = X$.

$$I_0 = \{x \in X_0\}$$

$$I_k^p(I_{k-1}^c, u_{c,k-1}) = \bigcup_{x\in I_{k-1}^c, u_{uc}\in U_{uc}, \delta\in\Delta} x + u + \delta \quad (10)$$

$$I_k^c = I_k^p \cap \hat{I}_k \quad (11)$$

Using the sets defined above and the set of stealthy attacks (8), we design an estimator system that receives as inputs the measurements, the control inputs, and the detector's statistics and computes the smallest state set that contains the actual state.

Definition 1: (*Estimator System*): Estimator system corresponding to system (1) employed with the detector (7) is

$$I_k(I_k^p, I_k^s, \tilde{x}_k) = \begin{cases} I_k^c & \tilde{x}_k \in I_k^s \\ \text{undefined} & \text{else} \end{cases} \quad (12)$$

Note that the set of stealthy attacks I_k^s describes how the detector (7) affects the estimator system. In particular, if a measurement $\tilde{x}$ is not in the set of stealthy attacks, i.e., $\tilde{x}_k \notin I_k^s$, the output of the estimator system is undefined because a detection alert is triggered. On the other hand, if $\tilde{x}$ belongs to the set of stealthy attacks, i.e., $\tilde{x}_k \in I_k^s$, the output of the estimator is the corrected set I_k^c defined by (11).

Remark: There is a tradeoff between detector's threshold and estimator performance. More specifically, if the detector's threshold increases, the size of stealthy attack space increases as well (Lemma 2), which leads to a larger set of estimated states (Equation (3)). Consequently, more restriction needs to be enforced by the supervisory controller to handle this larger set of states.

B. Supervisory Control Problem

To formulate the supervisory control problem, we first transform the infinite state space of the system model to a finite set of discrete states. Similar to [10], we define a set of discrete states Q and a mapping $\ell : X \to Q$ as follows:

$$\ell_i(x_i) := \begin{cases} c\tau\mu, \text{ for } c \in \mathbb{Z} \text{ s.t.} & \\ c\tau\mu - \tau\mu/2 < x_i \leq c\tau\mu + \tau\mu/2 & x_i \leq \beta_l \\ q_{i,m} & x_i > \beta_l \end{cases} \quad (13)$$

where l is the index of vehicle i's road ($i \in \Pi_l$). We define $\ell(x)$ as the vector $(\ell(x_1), ..., \ell(x_n))$. Further, we assume that for all $q \in Q$, there exists some $x \in X$ such that $\ell(x) = q$.

Let $2^X/\ell$ denote the quotient set of 2^X by the equivalence classes induced by ℓ. The supervisory control problem is stated below:

Problem 1: Given $2^X/\ell$, design a supervisor $\sigma : 2^X/\ell \to 2^{V_c}$ that associates to each $I_k \in 2^X$ a set of inputs $v_c \in V_c$ allowed for the interval $[k\tau, (k+1)\tau]$ with the following properties: *(i)* if $v_c(t) \in \sigma(I_k)$, then the trajectory of the system is safe in the interval $[k\tau, (k+1)\tau]$. *(ii)* if $v_c \in \sigma(I_k)$ and $\ell(I_{k+1}) \neq q_m$, then $\sigma(I_{k+1}) \neq \emptyset$. *(iii)* if $\tilde{\sigma} \neq \sigma_R$ and $\tilde{\sigma}$ is non-blocking and safe, then $\tilde{\sigma}(I_k) \subseteq \sigma_R(I_k)$.

C. Resilient Supervisor Design

We construct a DES abstraction of the estimator system and solve the supervisory control problem using the abstracted model. More specifically, we take the following steps: (1) Defining a DES with a set of events corresponding to controlled inputs, uncontrolled inputs, disturbance, prediction-correction, and detection such that the uncontrolled inputs and the disturbance events are unobservable. (2) Constructing an observer of the DES and providing conditions under which the observer is an abstraction of the estimator system (3) Translating safety and non-blockingness specifications to the DES domain, solving the supervisory control problem given the observer, and translating the obtained supervisor to the continuous time domain.

1) DES Model: We define a DES with an event set that models controlled inputs, uncontrolled inputs, disturbance, prediction-correction, and detection. This is similar to the approaches used in [10] and [11]. Consider a five-layer event set $E = \Lambda^d \times \Lambda^c \times U_c \times U_{uc} \times W$. The event sets U_c and U_{uc} contain inputs of controlled and uncontrolled vehicles and are defined as described in Section II. The set W represents a discretization of the set of disturbances Δ, defined by $W = \{k\tau\mu : \lfloor \delta_{min}/(\tau\mu) \rfloor \leq k \leq \lceil \delta_{max}/(\tau\mu) \rceil\}$. The sets Λ^c and Λ^d, which respectively represent prediction-correction and detection, are as defined below.

$\mathbf{\Lambda^c}$: Let any two inputs of $\hat{I}(\cdot)$ be equivalent, if they result in the same discrete set with respect to $\ell(\hat{I}(\cdot))$. That is to say, given (9) and (13), define the equivalence relation $(x_1, u_1 ... u_{T-1}, x_T) \equiv (x'_1, u'_1 ... u'_{T-1}, x'_T) \Leftrightarrow$
$\ell(\hat{I}(x_1, u_1 ... u_{T-1}, x_T)) = \ell(\hat{I}(x'_1, u'_1 ... u'_{T-1}, x'_T))$. Define $\phi := (x_1, u_1 ... u_{T-1}, x_T)$ and $\Phi := X \times U^T \times X$. For any $\phi \in \Phi$, let $[\phi]$ denote the equivalence class of ϕ under relation $\equiv$ and let Λ^c represent the set of equivalence classes. These equivalence classes form a finite set of discrete events $\lambda^c \in \Lambda^c$.

We say upon occurrence of event λ^c, there exists a transition from a state q to itself if and only if there exist some continuous element ϕ whose equivalence class is λ^c [11]. To represent this, define the discrete function $I^{\Lambda^c}(\lambda^c) : \Lambda^c \to 2^Q$ as

$$I^{\Lambda^c}(\lambda^c) = \{q \in Q : (\exists \phi \in \Phi : [\phi] = \lambda^c)[\hat{I}(\phi) \cap \ell^{-1}(q) \neq \emptyset]\}$$

Then, define the transition function $\psi^c : Q \times \Lambda^c \to Q$, which is shown to be the analogue of the estimate correction (11) in the DES domain, as

$$\psi^c(q, \lambda^c) = \begin{cases} q & \text{if } q \in I^{\Lambda^c}(\lambda^c) \\ \text{undefined} & \text{else} \end{cases}$$

$\mathbf{\Lambda^d}$: Let any two elements (x, u, C) and (x', u', C') be in the same equivalence class if their sets of discrete states with respect to $\ell(I^s(\cdot))$ are the same. That is, given (8), let $(x, u, C) \equiv (x', u', C') \Leftrightarrow \ell(I^s(x, u, C)) = \ell(I^s(x', u', C'))$. Define $\theta := (x, u, C)$, and $\Theta := X \times U \times \mathbb{R}$. For any $\theta \in \Theta$, let $[\theta]$ denote the equivalence class of θ under this relation. Denote by Λ^s the set of such equivalence classes. The finite set Λ^d is defined by $\Lambda^d := \Lambda^s \times Q$.

To define the transition function corresponding to Λ^d, define the function $I^{\Lambda^d} : \Lambda^s \times Q \to \{H_0, H_1\}$, for any $\tilde{q} \in Q$, by

$$I^{\Lambda^d}(\lambda^s, \tilde{q}) = \begin{cases} H_0 & \exists \theta \in \Theta : [\theta] = \lambda^s, I^s(\theta) \cap \ell^{-1}(\tilde{q}) \neq \emptyset \\ H_1 & \text{else} \end{cases}$$

The transition function $\psi^d : Q \times \Lambda^d \to Q$, which corresponds to the detector in the DES domain, is defined as

$$\psi^d(q, \lambda^d) = \begin{cases} q & \text{if } I^{\Lambda^d}(\lambda^d) = H_0 \\ \text{undefined} & \text{else} \end{cases}$$

where $I^{\Lambda^d}(\lambda^d) \neq H_0$ represents detection of attacks.

Define the DES

$$\tilde{G} := (\tilde{Q}, E, \psi, q_0, Q_m) \tag{14}$$

with state set $\tilde{Q}$, event set E, transition function $\psi : Q \times E \to Q$, set of initial states $Q_0 \subset Q$, and set of marked states $Q_m = \{q_m\}$. Transition function ψ has the form $\psi(q, \lambda^d, \lambda^c, u_c, u_{uc}, w) =$
$\psi_3(\psi_2(\psi_1(\psi^c(\psi^d(q, \lambda^d), \lambda^c), u_c), u_{uc}), w)$, and the language of this system is $(\Lambda^d \Lambda^c U_c U_{uc} W)^*$. To avoid state conflicts, we define intermediate states Q', Q'', Q_{I1}, and Q_{I2} such that $\psi^d : Q \times \Lambda^d \to Q'$, $\psi^c : Q' \times \Lambda^c \to Q''$, $\psi_1 : Q'' \times U_c \to Q_{I1}$, $\psi_2 : Q_{I1} \times U_{uc} \to Q_{I2}$, and $\psi_3 : Q_{I2} \times W \to Q$. The uncontrolled inputs U_{uc} and the disturbances W are not directly observed so we define them as uncontrollable and unobservable. Further, we let the controllable inputs U_c be observable and controllable, and Λ^c and Λ^d be observable but uncontrollable [11].

2) Observer: System $\tilde{G}$ is a discrete event system with unobservable events and we can design its observer $Obs(\tilde{G})$ [7]. The states of the observer are information states $\iota \subseteq 2^{\tilde{Q}}$ and its transition functions are: $\bar{\psi}^d : 2^Q \times \Lambda^d \to 2^{Q'}$, $\bar{\psi}^c : 2^{Q'} \times \Lambda^c \to 2^{Q''}$, and $\bar{\psi}^p : 2^{Q''} \times U_c \to 2^Q$. Note that by construction, all controllable events of $Obs(\tilde{G})$ are observable and thus, an optimal supervisor solution indeed exists.

We show that if δ_{min} and δ_{max} are multiples of μ, the DES observer automaton constructed above is an abstraction of (12) and its controller can be used to control the continuous system. We use the notions of simulation and alternating simulation [19]. In what follows, consider the observation

maps $H_{2^X}(I) = \ell(I)$ and $H_{2^Q}(\iota) = \iota$, and the relation $R = \{(\iota, I) \in 2^Q \times 2^X : \iota = \ell(I)\}$.

Lemma 4: If δ_{min} and δ_{max} are multiples of μ, system (12) alternatingly simulates $Obs(\tilde{G})$, and $Obs(\tilde{G})$ simulates system (12).

Proof: (System (12) alternatingly simulates $Obs(\tilde{G})$): 1) By assumption, for every $\iota_0 \subseteq 2^{Q_0}$ there exists $I_0 \subseteq 2_0^X$ with $\ell(I_0) = \iota_0$; 2) By definition of the observation maps, for every $\ell(I) = \iota$ we have $H_{2^Q}(\iota) = H_{2^X}(I)$; 3) We need to prove for every $\iota = \ell(I)$ and for every $u \in U_{2^Q}(\iota)$ there exists $u' \in U_{2^X}(I)$ such that for every (I, u, I') there exists (ι, u', ι') satisfying $\ell(I') = \iota'$. Consider any transition (I, u, I') and let $u = u'$. The sets I and ι can be written as $I = [x_l, x_h]$ and $\iota = \ell(I) = \{[q_l, q_h]\}$. Because δ_{min} and δ_{max} are multiples of μ, we have $\bar{\psi}^p(\iota, u) = \ell(I^p)$ [9]. Then, $\iota' = \bar{\psi}^c(\bar{\psi}^d(\bar{\psi}^p)) = \ell(I^p) \cap \ell(\hat{I}) = \ell(I^p \cap \hat{I}) = \ell(I')$.

($Obs(\tilde{G})$ simulates system (12)): 1) By assumption, for every $I_0 \subseteq 2^X$ there exists $\iota_0 \subseteq 2^Q$ with $\ell(I_0) = \iota_0$; 2) It is similar to alternating simulation; 3) Let $u = u'$. We know if $\ell(I) = \iota$, then $\ell(I^p) = \bar{\psi}^p(\iota, u_c)$. By definition, $\ell(I') = \ell(I^p \cap \hat{I}(\phi))$ and by construction there exist $\lambda^c = [\phi]$ and $\lambda^s = [\theta]$ such that $\ell(I^p \cap \hat{I}) = \bar{\psi}^c(\bar{\psi}^d(\bar{\psi}^p, \lambda^d), \lambda^c) = \iota'$. Hence, for every (I, u, I') there exists (ι, u, ι') such that $\ell(I') = \iota'$. ∎

3) Supervisor Solution: To construct a safe, non-blocking, and maximally permissive supervisor at the DES level, we use the solution to the Basic Supervisory Control Problem in the Non-Blocking case (BSCP-NB) as described in [17], [7]. The solution computes the supremal controllable sublanguage of a specification with respect to the language of a system. More specifically, it computes the language $(\mathcal{L}_m(H))^{\uparrow C}$, where H is a specification and $\uparrow C$ is the supremal controllable sublanguage operation [13]. The BSCP-NB algorithm constructs a supervisor S such that $\mathcal{L}_m(S/G) = (\mathcal{L}_m(H))^{\uparrow C}$ and $\mathcal{L}(S/G) = \overline{(\mathcal{L}_m(H))^{\uparrow C}}$, where S/G is the system G controlled by S and $\overline{L}$ denotes the prefix closure of language L [7].

We define the transition function, safety specification, and set of marked states for the estimator system and then translate it to DES domain [11]. The estimator system (12) defines a transition system as follows:

$$(I, u, I') \in \rightarrow \subseteq 2^X \times U_c \times 2^X \text{ if } \exists \tilde{x} \in X, I(I^p, I^s, \tilde{x}) = I'$$

Define a transition from (I, u, I') as safe if all trajectories of the continuous time system corresponding to this transition are safe:

$$(I, u, I') \in Safe \subseteq \rightarrow \text{ if } \forall x \in I, x' \in I', (x, u, x') \in \rightarrow_{(2)} \Rightarrow (x, u, x') \in Safe_{(2)} \quad (15)$$

Further, define the set of marked states as:

$$I_m := \{I \subseteq 2^X : \ell(I) = q_m\} \quad (16)$$

The safety specification (15) and the set of marked states (16) are translated to the DES domain as follows:

$$(\iota, u, \iota') \in Safe_{\tilde{G}} \subseteq 2^Q \times U \times 2^Q \text{ if } \forall (I, u, I') \in \rightarrow, \iota = \ell(I), \iota' = \ell(I') : (I, u, I') \in Safe \quad (17)$$

$$\iota_m := \{\{q_m\}\} \quad (18)$$

Solving BSCP-NB for $Obs(\tilde{G})$ with respect to the specifications (17) and (18), outputs a safe, non-blocking, and maximally permissive supervisor [7]. Denote by S this safe, non-blocking, and maximally permissive supervisor. Define the supervisor map

$$\sigma(I) = \{u_c/\tau : u_c \in S(\ell(I))\} \quad (19)$$

We call (19), the resilient supervisory controller of system (1).

Theorem 1: The supervisory controller (19) that associates a set of admissible inputs $v_k \in V_c$ to each $x_k \in X$, allowed for the interval $t = [k\tau, (k+1)\tau]$ solves Probem 1.

V. Example

We consider two controllable vehicles driving on separate roads approaching an intersection. Consider the set of inputs $V_c = \{1, 3\}$, disturbances $d = [0, 1]^2$, time-discretization $\tau = 1$, space discretization parameter $\mu = 1$, maximum attack length $T_{max} = 1$, and detector threshold $\eta \to 0$. Let $x_0 = (1, 1)$, and suppose $(\alpha_1, \beta_1) = (\alpha_2, \beta_2) = [9.5, 12.5]$, so the vehicles collide if at a time t, $9.5 \leq x_1(t), x_2(t) \leq 12.5$. A marked state is reached if $x > 12.5$ for both vehicles.

We define an attack model known as surge attack. In such attacks, the adversary adds the maximum error value that bypasses the detector, that is to say, he keeps the detector's statistic at the threshold without exceeding it [5]. This required solving the equation $C(k) + z(k) - b = \eta$, for every $k \in [k_s, k_e)$. Using Lemma 2, the resulting attack at $k = k_s$ is

$$\tilde{x}_{k_s} = \hat{x}_{max,k_s} + \eta + b$$

where $\hat{x}_{max,k_s} = x_{k_s-1} + u_{k_s-1} + d_{max}$. For the remainder of the attack, which is $k \in [k_s + 1, k_e)$, we have

$$\tilde{x}_k = \tilde{x}_{k-1} + u_{k-1} + d_{max} + b$$

Note that we can replace $\hat{x}_{max}$ and d_{max} with $\hat{x}_{min}$ and d_{min} to achieve similar results for negative error values.

To show how sensor attacks compromise safety, we first consider the supervisor designed in [10] as the controller of the intersection. Constructing the supervisor, the admissible inputs for each cell $q \in Q$ are shown in Fig. 5. Note that since $n = 2$, the bad set B is bounded. The set of enabled inputs for the initial state is $\tilde{\sigma}(x(0)) = \{(1, 3), (3, 1)\}$. Suppose the vehicles select $u = (1, 3)$ and the state evolves to $x(1) = x(0) + u(0) + (0, 0) = (2, 4)$. The attacker receives the actual measurement $x(1)$ and starts a surge attack on vehicle 2 by transmitting the corrupted vector $\tilde{x}(1) = x(1) + (0, 1) = (2, 5)$. The detector will not trigger an alert because $\tilde{x}(1) \in I_1^s$. The set of admissible inputs for the

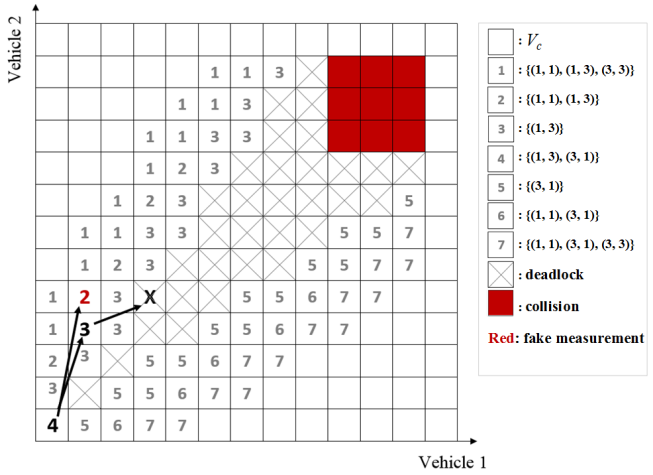

Fig. 5: Example of a stealthy sensor attack that compromises safety.

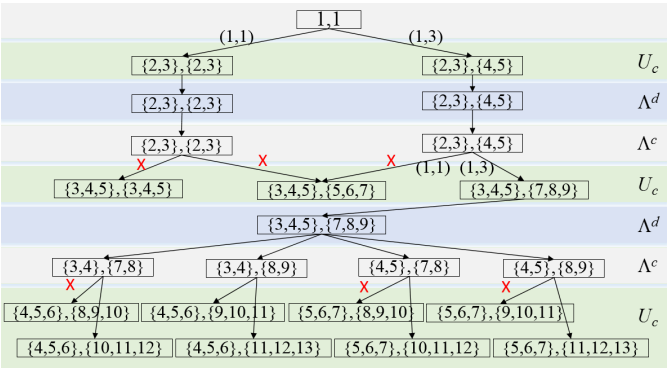

Fig. 6: Observer of the two-vehicle intersection.

corrupted state is $\tilde{\sigma}(\tilde{x}(1)) = \{(1,1),(1,3)\}$, whereas for the actual state this set is $\tilde{\sigma}(x(1)) = \{(1,3)\}$. The input $(1,1)$ is incorrectly enabled by the supervisor in the perceived state and hence, selecting it may result in deadlock (i.e., unavoidable collision).

Next, we construct the resilient supervisor for the example and show how the discussed stealthy attack is handled by the supervisor. To implement the resilient supervisor, we need to construct the observer of system $\tilde{G}$ as shown in Fig. 6. The figure only shows the different choices of inputs for vehicle 2 and input of vehicle 1 is set to $v_1 = 1$. Similar to the previous case, let $u(0) = (1,3)$ be the first chosen input so the information state of the system becomes $(\{2,3\}\{4,5\})$. Assume the actual state is $x(1) = (2,4)$ and the attacker transmits the corrupted measurement $\tilde{x}(1) = (2,5)$. The events $\lambda^c = [\tilde{x}(0), u(0), \tilde{x}(1)]$ and $\lambda^d = [\tilde{x}(0), u(0), C(0)]$ occur and the information state transits to $\iota(1) = (\{2,3\}\{4,5\})$. The input $(1,1)$, which caused collision in the previous case, is disabled in this scenario since if selected, the vehicles may enter any of the illegal states $(4,5),(5,5)$, and $(5,6)$. Therefore, another input is selected (e.g., $u(1) = (1,3)$) and the vehicles remain safe despite the attack.

To show how the state evolves in the following steps,

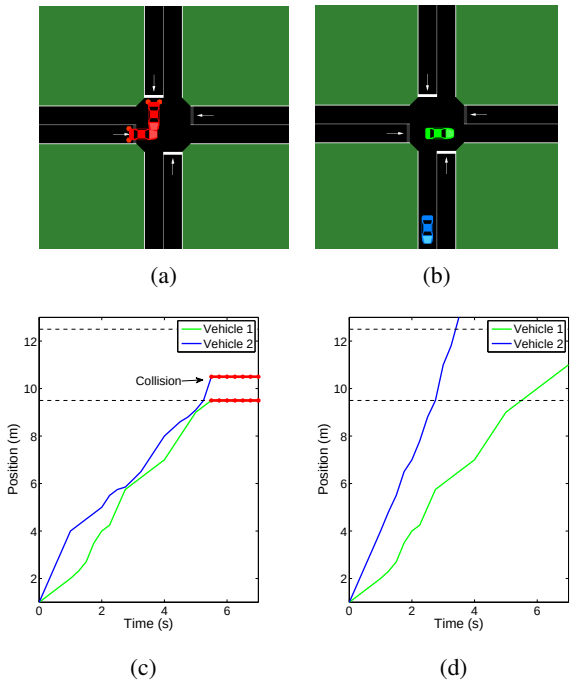

Fig. 7: Simulation of the example in SUMO. Figures 7a and 7c correspond to the non-resilient supervisor and Figures 7b and 7d correspond to the resilient supervisor. The area between the dashed lines represents the intersection.

we construct the observer for an additional sequence of events as shown in Fig. 6. Assume the current information state is $(\{3,4,5\},\{7,8,9\})$. Upon occurrence of an event λ^c, this information state has 4 possible successors. For instance, if the new measurement is $\tilde{x}(2) = (3,8)$, then $\lambda^c = [(\tilde{x}(1), u(1), \tilde{x}(2))]$ with $\tilde{x}(1) = (2,5)$, $u(1) = (1,3)$, and $\tilde{x}(2) = (3,8)$, and the successor state is $\iota(2) = (\{3,4\},\{8,9\})$. This shows how the measurements $\tilde{x}(1)$ and $\tilde{x}(2)$, and the input $u(1)$ are used to correct the predicted set of states to a subset $\iota(2)$.

To simulate the scenarios, we use Simulation of Urban MObility (SUMO). SUMO is an open source platform for microscopic road traffic simulation [15]. The objects of the SUMO simulation can be controlled through Traffic Control Interface (TraCI), which uses a TCP-IP based client/server architecture. We use the TraCI4Matlab implementation of TraCI.

The simulations and their corresponding trajectories are shown in Fig. 7. We assume vehicle 1 drives the same way in both scenarios. Also, we assume following an unavoidable collision, vehicles keep driving with the same speed. The simulations confirm our previous discussion that the resilient supervisor maintains safety even in the presence of sensor attacks. In this case, this is done by disabling the unsafe input $(1,1)$ at $t = 1$, and enabling $(1,3)$ at subsequent steps.

Although safety can be maintained for any finite T_{max}, the higher T_{max}, the more restrictive the supervisor becomes. In other words, there exists a tradeoff between T_{max} and supervisor's performance. To demonstrate this,

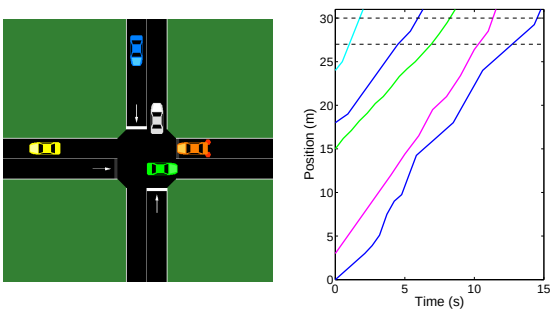

Fig. 8: System remains safe if every pair of vehicles is safe.

one can extend the example by considering different values of T_{max}. If $T_{max} \geq 4$, inputs $\{(1,3),(3,1)\}$ are enabled at the initial state. Then, until a marked state is reached, the supervisor disables all the inputs except the one chosen at the initial state. This is clearly more restrictive than the case of $T_{max} = 1$. Nevertheless, for this example, even if $T_{max} = (x_m - x_0)/(v_{min} + d_{min}) = 12$, a solution exists.

Scenarios including multiple vehicles can be similarly analyzed by verifying safety for each pair of vehicles (as shown in Fig. 8 for a different network with $n = 5$). The scalability of the proposed resilient supervisory control system depends on the number of vehicles, number of control inputs, minimum and maximum disturbances, discretization parameters, and maximum attack length. Scalability can be characterized based on the computational complexity for implementing supervisors for partially-observed DES [21].

VI. Conclusion

We studied the problem of supervisory control of autonomous intersections in the presence of sensor attacks. We showed that the supervisory control system is vulnerable to sensor attacks that can cause collision or deadlock among vehicles. To improve the system resilience, we introduced a detector in the control architecture with the purpose of detecting attacks. There exist stealthy attacks that cannot be detected but are capable of compromising safety. To address this issue, we designed a resilient supervisor that maintains safety even in the presence of stealthy attacks. The resilient supervisor consists of an estimator that computes the smallest state estimate compatible with the control inputs and measurements seen thus far. We formulated the resilient supervisory control system and presented a solution. We demonstrated how the resilient supervisor works by considering examples.

Our solution can be extended to handle other attack models. Infinite length attacks and DoS attacks are special cases of the attack model discussed in this work and can be studied in a similar manner. Even though we focused on the analysis of supervisory control of autonomous intersections, our techniques can be applied to other cyber-physical system as well. Future work include study of actuator attacks, development of efficient methods for implementing resilient controllers, and construction of decentralized resilient controllers.

References

[1] H.-S. Ahn, A. Colombo, and D. Del Vecchio. Supervisory control for intersection collision avoidance in the presence of uncontrolled vehicles. In *American Control Conference (ACC), 2014*, pages 867–873. IEEE, 2014.

[2] S. Amin, A. A. Cárdenas, and S. S. Sastry. Safe and secure networked control systems under denial-of-service attacks. In *Hybrid Systems: Computation and Control*, pages 31–45. Springer, 2009.

[3] S. Azimi, G. Bhatia, R. Rajkumar, and P. Mudalige. Reliable intersection protocols using vehicular networks. In *Cyber-Physical Systems (ICCPS), 2013 ACM/IEEE International Conference on*, pages 1–10. IEEE, 2013.

[4] M. Basseville and I. V. Nikiforov. *Detection of Abrupt Changes: Theory and Applications*. 1993.

[5] A. A. Cárdenas, S. Amin, Z.-S. Lin, Y.-L. Huang, C.-Y. Huang, and S. Sastry. Attacks against process control systems: Risk assessment, detection, and response. In *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security*, ASIACCS '11, pages 355–366, New York, NY, USA, 2011. ACM.

[6] A. A. Cárdenas, S. Amin, and S. Sastry. Secure control: Towards survivable cyber-physical systems. In *The 28th International Conference on Distributed Computing Systems Workshops*, pages 495–500. IEEE, 2008.

[7] C. G. Cassandras et al. *Introduction to discrete event systems*. Springer Science & Business Media, 2008.

[8] A. Colombo and D. Del Vecchio. Supervisory control of differentially flat systems based on abstraction. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 6134–6139. IEEE, 2011.

[9] E. Dallal. *Discrete event system methods for control problems arising in cyber-physical systems*. PhD thesis, The University of Michigan, 2014.

[10] E. Dallal, A. Colombo, D. Del Vecchio, and S. Lafortune. Supervisory control for collision avoidance in vehicular networks using discrete event abstractions. In *American Control Conference (ACC), 2013*, pages 4380–4386, June 2013.

[11] E. Dallal, A. Colombo, D. Del Vecchio, and S. Lafortune. Supervisory control for collision avoidance in vehicular networks with imperfect measurements. In *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, pages 6298–6303. IEEE, 2013.

[12] K. Dresner and P. Stone. A multiagent approach to autonomous intersection management. *Journal of artificial intelligence research*, pages 591–656, 2008.

[13] X. D. Koutsoukos, P. J. Antsaklis, J. A. Stiver, and M. D. Lemmon. Supervisory control of hybrid systems. *Proceedings of the IEEE*, 88(7):1026–1049, 2000.

[14] H. Kowshik, D. Caveney, and P. Kumar. Provable systemwide safety in intelligent intersections. *Vehicular Technology, IEEE Transactions on*, 60(3):804–818, 2011.

[15] D. Krajzewicz, J. Erdmann, M. Behrisch, and L. Bieker. Recent development and applications of sumo–simulation of urban mobility. *International Journal On Advances in Systems and Measurements*, 5(3&4), 2012.

[16] F. Pasqualetti, F. Dorfler, and F. Bullo. Attack detection and identification in cyber-physical systems. *Automatic Control, IEEE Transactions on*, 58(11):2715–2729, 2013.

[17] P. J. Ramadge and W. M. Wonham. Supervisory control of a class of discrete event processes. *SIAM journal on control and optimization*, 25(1):206–230, 1987.

[18] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. C. Teneketzis. Failure diagnosis using discrete-event models. *Control Systems Technology, IEEE Transactions on*, 4(2):105–124, 1996.

[19] P. Tabuada. *Verification and control of hybrid systems: a symbolic approach*. Springer Science & Business Media, 2009.

[20] Q. Wen, R. Kumar, J. Huang, and H. Liu. A framework for fault-tolerant control of discrete event systems. *Automatic Control, IEEE Transactions on*, 53(8):1839–1849, 2008.

[21] X. Yin and S. Lafortune. A general approach for synthesis of supervisors for partially-observed discrete-event systems. *19th IFAC World Congress*, pages 2422–2428, 2014.

[22] F. Zhao, X. Koutsoukos, H. Haussecker, J. Reich, and P. Cheung. Monitoring and fault diagnosis of hybrid systems. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 35(6):1225–1240, 2005.